

ADAPTABLE CACHE FOR DYNAMIC DIGITAL MEDIA

FIELD OF THE INVENTION

[0001] This invention relates to the field of digital media servers.

BACKGROUND OF THE INVENTION

5 **[0002]** Digital media services are growing in importance and gaining acceptance as the underlying technologies for providing these services mature. But, one limitation on implementing such services is the transactional performance or data-access speed of storage devices. As the demand for digital media increases, limits on transactional performance become more evident and impede system performance and growth.

10 **[0003]** One known technology for improving data-access performance is a cache. A cache is high-speed memory accessible by a disk, controller, or host, whose primary purpose is to reduce the effective time required to read data from or write data to a lower speed storage device. A cache may reside within a storage device, host processor, or controller, or may be located external to any of these devices.

15 **[0004]** Caches improve performance in hard disk based systems by reducing the requirement for mechanical movement in the storage device, thus minimizing or eliminating bottlenecks related to hard disk access speed. Such caches typically utilize random access memory (RAM) to speed up data access by temporarily storing frequently or most-recently used data in the cache. Although RAM access is much faster than hard disk access, cache effectiveness is often hindered by latencies associated with multiple bus traversals and the fixed amount of RAM that may be integrated into the cache. Furthermore, hard disk drive cache sizes are usually limited to 8MB or less. These caches lack the physical size to store requested resources in the media server context, where a typical resource is in excess of 100MB.

25 **[0005]** Caches implemented on other devices, such as a processor or host bus adapter, typically utilize a small amount (512 KB or less) of high speed memory for caching frequently used instructions or data. These caches are typically too small to increase system performance where large amounts of data are being transferred between devices or system buses. Furthermore, system performance is typically still limited, since points of system

degradation ("bottlenecks") may still occur at connections between devices and I/O buses or at connections between the cache and the device accessing the cache.

[0006] Fig. 1 illustrates a typical media server arrangement in the prior art. As shown in Fig. 1, media server 100 typically comprises a host processor 120 which is a computing device having a central processing unit adapted to access and process instructions stored in a memory 122. Host processor 120 preferably runs an application adapted to receive digital media requests from clients, administrators, and other applications (e.g., an application running on a second media server in the same server cluster as media server 100), and to respond to such requests by retrieving the requested digital media, converting it to wire format, and transmitting it to a client or other location via a suitable network or other communication link.

[0007] Media server 100 further comprises a network interface 130 for receiving digital-media assets asset requests (e.g., for a movie, video clip, etc.) and forwarding the requests to host processor 120 via a connection 146B. Network interface 130 may, for example, be a network interface card (NIC). Connection 146B connects to I/O bus 106B which transmits the request to host processor 120 via connection 148A. Host processor 120 processes received requests and retrieves requested assets from a storage system 102 via host bus 106B and a secondary (storage system) I/O bus 106A, as described below.

[0008] Also shown in Fig. 1 is a storage system 102 comprising a plurality of interconnected physical or logical storage devices 104. Each storage device 104 may contain one or more hard disk drives 180.

[0009] Further shown in Fig. 1 is a controller 128 preferably adapted to facilitate communication between host I/O bus 106B and storage system I/O bus 106A. Host I/O bus 106B preferably is connected to network interface 130 via bi-directional connection 146, and to host processor 120 via bi-directional connections 148. Similarly, storage system I/O bus 106A preferably is connected to storage devices 104 via bi-directional connections 140.

[00010] Figs. 2A-2D are composite block/flow diagrams illustrating processing of a typical asset request in the prior art. As shown in Fig. 2A, a user request preferably is received at network interface 130 and forwarded to host processor 120 via connection 146B, I/O bus 106B, and connection 148A. In Fig. 2B, host processor 120 processes the request and sends instructions to retrieve the requested asset through host side I/O bus 106B, controller 128 (e.g., a host bus adapter), and storage system I/O bus 106A to the storage device 104 that contains the asset. In some systems, storage system 102 may reside on the same I/O bus as host processor 120.

[00011] In Fig. 2C, storage device 104 retrieves the requested asset and delivers it to host processor 120. Host processor 120 prepares the asset and transmits it to network interface 130 for delivery to the requesting user as shown in Fig. 2D.

[00012] Fig. 3 is a block diagram illustrating potential bottlenecks in media server 100.

5 As shown in Fig. 3, bottlenecks may be device related or may occur at connections between devices and I/O buses, at connections between two or more I/O buses (e.g., at controller 128), or as a result of the physical limitations of the I/O bus itself. These bottlenecks decrease speed and efficiency of the overall system by requiring components to operate at less than their optimal individual speeds.

10 [00013] More specifically, asset requests received at network interface 130 may overburden the interface and result in bottleneck 360. Further, the media server may encounter system performance degradation resulting in bottleneck 350 if host processor 120 cannot process requests efficiently. Additionally, if storage device 104 cannot retrieve the requested asset immediately, or if storage device 104 is limited by its physical characteristics,
15 bottleneck 310 may result.

[00014] In addition to these device-related bottlenecks, connection-related bottlenecks may also occur. For example, where the connection between storage device 104 and the storage system I/O bus has a lower bandwidth capability than the storage device or the I/O bus, bottleneck 320 may occur. As described below, this results in storage device 104
20 restricting its data flow to I/O bus 106 due to the bandwidth capability of connection 140 and not utilizing its full capability for transmitting information to storage system I/O bus 106A. Further, where controller 128 cannot route data efficiently from one I/O bus to another, or where controller 128 becomes overburdened due to bus traversal traffic, bottleneck 330 may occur. Additionally, when the I/O bus itself cannot handle the amount of data passed to it
25 from a device, bottleneck 340 may occur. As will be recognized, bottlenecks may also occur at any inter-connection between devices where data flow exceeds the connection bandwidth, e.g., where data transmissions from I/O bus 106A to storage device 104 exceed the maximum bandwidth capability of the connection interface.

[00015] Fig. 4 illustrates potential bottleneck 320 in greater detail. As shown in Fig. 4,
30 storage device 104 may have a bandwidth capability 410 comparable to the bandwidth capability 420 of I/O bus 106. By contrast, connection 140 may be a typical interface connection with a bandwidth capability 430 of less than 50 MBps, such as IDE (ATA), or E-IDE (fast ATA). Thus, while the storage device may be capable of delivering enough data to saturate the I/O bus, interface connection 140 nevertheless restricts the flow of data. As will

be recognized, this bottleneck may be alleviated to some degree by using an alternative interface such as an optical interface (fibre channel arbitrated loop ("FC-AL")), small computer system interface ("SCSI"), or serial ATA ("S-ATA") connection, but some degree of bottleneck may nevertheless remain.

5 **[00016]** Figs. 5A-5B are a block diagram illustrating the number of input and output operations per second (IOPS) a storage device typically performs. The higher the number of IOPS a device can perform, the better the resulting performance of the particular device and, in general, the entire system. While hard disk drives have been increasing in speed and capacity, inherent physical limitations associated with retrieving data from the rotating
10 platters still exist.

[00017] One aspect of these physical limitations is measurable as seek times, i.e., the amount of time it takes for the drive to access the beginning of a block of data to be transferred from the disk, given the starting and target locations. Seek times on the highest performing disk drives today average in the 3-5 millisecond (ms) range for randomly located
15 data. Assuming a seek time of 5 ms (and without taking into consideration the time needed to read from the drive) the drive would theoretically be limited to between 200 IOPS. A faster seek time of 3 ms would result in a maximum IOPS of 333.

[00018] In addition to seek time, however, another important aspect that reduces IOPS is transfer time. Most high performance drives can sustain transfers of approximately 25MB
20 per second. This number is an average since blocks on the inner portion of the platter can be accessed faster than blocks on the outer portion of the platter. For example, if a media server needs to transfer one second of content for a movie encoded at 3.75Mbps, approximately 3.75 Mb (or 0.469 MB) of data would be transferred. This represents approximately 1/50th of a second or approximately 20ms. Coupled with the initial seek time (3-5 ms) from above,
25 typical time requirements are thus approximately 23-25ms. Using this approximation, a typical media server with only one drive can perform between 40-43 IOPS.

[00019] A primary method of increasing IOPS performance, often incorporated into high-performance media servers and storage systems, is to utilize a plurality of directly connected hard disk drives. This can improve performance by performing transactions in
30 parallel on different drives. It also can improve performance by decreasing the transfer time. Generally, high-performance hard disk drives are used for increased system performance. A high-performance hard disk drive typically has a maximum mechanical spindle rotation speed of 15,000 revolutions per minute (RPM). In a digital media server or stand-alone storage system, these hard disk drives are either configured as just a bunch of disks (JBOD) or, more

typically, as a redundant array of independent disks (RAID). Using one of these configurations increases the virtual size of the storage system, allowing a larger amount of data to flow from the virtual disk. For example, four hard disk drives spinning at 15,000RPM can be represented as a single virtual disk, which theoretically achieves a maximum rotational speed of 60,000RPM.

[00020] As shown in Figs. 5A-B, if a single 15,000RPM hard disk drive may achieve IOPS performance of X, then four identical drives in a virtual set can theoretically achieve IOPS performance of 4X. For example, if the one drive delivers 40-43 IOPS, then the 4 drives can theoretically achieve 160-172 IOPS. As will be recognized, however, there are many real-world factors, such as data placement on the drive platters, file fragmentation, and head placement at the time of the request, that make this theoretical limit virtually impossible to achieve outside the most ideal controlled conditions.

[00021] Thus, digital media systems based solely hard disk drives suffer from limitations on transactional throughput which manifest themselves in the form of relatively low IOPS and a reduced ability to deliver high numbers of digital media files to an ever increasing consumer base.

SUMMARY OF THE INVENTION

[00022] A system and method are disclosed for eliminating many of the transactional performance limitations in current digital media server systems by augmenting those existing systems with an adaptable cache. In a preferred embodiment, the adaptable cache is a compact storage device that can persist data and deliver it at an accelerated rate, as well as act as an intelligent controller and director of that data. Incorporating such an adaptable cache between existing storage devices and an external network interface of a media server, or at the network interface itself, significantly overcomes the transactional limitations of the storage devices, increasing performance and throughput for the overall digital media system.

[00023] The adaptable cache of the present system and method may preferably be integrated directly into the storage and delivery pipelines, utilizing the native communications busses and protocols of those subsystems. Further, the invention can dynamically incorporate algorithms for maximization of storage retrieval and delivery performance.

[00024] The present adaptable cache decreases the number of storage devices required by a digital media system to achieve a desired output capacity, thus reducing the physical and

administrative scale of the resulting system. This serves to increase the efficiency of all digital media servers in a system, reduce the total cost of system ownership, and make the system available to a larger number of consumers. Moreover, with the present system and method, system operators need not replicate digital media across large numbers of digital media servers or storage systems to overcome the transactional performance of inferior systems. This efficiency gain allows digital media system owners and operators to implement full-scale digital media services deployments as they become more cost effective to meet anticipated or actual user demand for those services.

[00025] In one aspect, the present invention is directed to a method for reducing bus traversal in a media server comprising a host processor, a network interface, and a storage subsystem comprising one or more storage devices, the host processor and network interface being connected to a first input-output bus, the storage subsystem being connected to a second input-output bus, the first and second input-output buses being connected via a controller, the method comprising:

providing an adaptable cache connected to the first input-output bus, said adaptable cache comprising a data interface, core logic, and electronic storage media;
receiving a request for a media asset via a network, said request being received by the network interface;

receiving the request at the adaptable cache;

processing the request by the adaptable cache, wherein if the requested media asset is found on the electronic storage media, the media asset is returned to the user via the first bus and not the second bus, and wherein if the requested media asset is not found on the electronic storage media, the media asset is accessed from the storage subsystem and returned to the user via the second bus and first bus.

[00026] In another aspect of the present invention, the method further comprises the request being received at the adaptable cache via the host processor.

[00027] In another aspect of the present invention, the method further comprises the request being received at the adaptable cache directly from the network interface.

[00028] In another aspect of the present invention, the method further comprises the adaptable cache being integrated with the network interface.

[00029] In another aspect of the present invention, the method further comprises the adaptable cache being integrated in the controller.

[00030] In another aspect of the present invention, the method further comprises the adaptable cache monitoring requests for media assets and if it is determined that the media

asset should be cached, the media asset is transferred from one or more storage devices to the electronic storage media.

[00031] In another aspect of the present invention, the method further comprises the adaptable cache monitoring requests for media assets and if it is determined that the media should be cached, the adaptable cache notifies requesting applications that it can accept future requests for said media assets.

[00032] In another aspect of the present invention, the method further comprises the adaptable cache monitoring requests for media assets and if it is determined that the media should be cached, the adaptable cache notifies the storage subsystem to disregard requests to deliver the media.

[00033] In another aspect of the present invention, the method further comprises wherein the requested media asset is not found on the electronic storage media, the adaptable cache stores the requested media asset on the electronic storage media.

[00034] In another aspect of the present invention, the method further comprises the adaptable cache integrating into the media server via an expansion card slot.

[00035] In another aspect of the present invention, the method further comprises the adaptable cache integrating with native communications busses and protocols existing on the media server.

[00036] In another aspect of the present invention, the method further comprises the adaptable cache utilizing the busses and protocols existing on the media server.

[00037] In another aspect, the present invention is directed to a method for improving transactional performance in a media server comprising a host processor, a network interface, and a storage subsystem comprising one or more storage devices, the host processor and network interface being connected to a first input-output bus, the storage subsystem being connected to a second input-output bus, the first and second input-output buses being connected via a controller, the method comprising:

providing an adaptable cache connected to the second input-output bus, said adaptable cache comprising a data interface, a core logic, and electronic storage media, and enabled to:

dynamically accept algorithms and heuristics that define or alter its operating characteristics without disrupting the operation of the media server,
determine whether to retrieve and store data from the storage subsystem based on the algorithms and heuristics;
retrieve data from the storage subsystem using its own data interface;

receiving a request for a media asset via a network, said request being received by the network interface;

receiving the request at the adaptable cache;

processing the request by the adaptable cache, wherein if the requested media asset is found on the electronic storage media, the media asset is returned to the user without accessing the one or more storage devices on the storage subsystem, and wherein if the requested media asset is not found on the electronic storage media, the media asset is accessed from one or more storage devices on the storage subsystem and returned to the user.

[00038] In another aspect of the present invention, the method further comprises the request being received at the adaptable cache via the second input-output bus.

[00039] In another aspect of the present invention, the method further comprises the adaptable cache integrating into the media server via an expansion card slot.

[00040] In another aspect of the present invention, the method further comprises the requested media asset not being found on the electronic storage media, and the adaptable cache then storing the requested media asset on the electronic storage media.

[00041] In another aspect of the present invention, the method further comprises the adaptable cache monitoring requests for media assets and if it is determined that the media asset should be cached, the media asset is transferred from one or more storage devices to the electronic storage media.

[00042] In another aspect of the present invention, the method further comprises the adaptable cache monitoring requests for media assets and if it determines that the media should be cached, the adaptable cache accepts future requests for said media assets.

[00043] In another aspect of the present invention, the method further comprises the adaptable cache monitoring requests for media assets and if it is determined that the media should be cached, the adaptable cache notifies the storage subsystem to disregard requests to deliver the media.

[00044] In another aspect of the present invention, the method further comprises wherein the adaptable cache integrates with native communications busses and protocols existing on the media server.

[00045] In another aspect of the present invention, the method further comprises the adaptable cache utilizing the busses and protocols existing on the media server.

[00046] In another aspect, the present invention is directed to a system for facilitating delivery of media resources, comprising:

a media server comprising a host processor, a network interface, and a storage

subsystem comprising electronic storage media, the host processor and network interface being connected to a first input-output bus, the storage subsystem being connected to a second input-output bus, the first and second input-output buses being connected via a controller,

5 an adaptable cache connected to an input-output bus comprising a data interface, core logic, and electronic storage media, the adaptable cache being adapted to store data on the electronic storage media, and further being adapted to receive and process requests for media assets, wherein if the requested media asset is found on the electronic storage media, the media asset is returned to a requestor via one or more I/O buses, and wherein if the requested
10 media asset is not found on the electronic storage media, the media asset is accessed from the storage subsystem and returned to the requestor.

[00047] In a preferred embodiment, the adaptable cache comprises core logic that includes instructions for when and how to cache content. This embeds decision making and control functionality directly into core logic, allowing the adaptable cache to become a hybrid
15 storage device/HBA. The core logic is also preferably programmed with dynamic algorithms and heuristics that define or alter the adaptable cache's operating characteristics.

BRIEF DESCRIPTION OF THE DRAWINGS

[00048] Fig. 1 is a composite block/flow diagram illustrating a typical media server of the prior art;

20 [00049] Figs. 2A-2D are block/flow diagrams illustrating data flow in a typical media server as known in the art;

[00050] Fig. 3 is a block/flow diagram illustrating potential bottlenecks in a typical media server of the prior art;

[00051] Fig. 4 is a block diagram illustrating one bottleneck of Fig. 3 in greater detail;

25 [00052] Figs. 5A-5B are block diagrams illustrating typical storage device performance in the prior art;

[00053] Fig. 6A is a block diagram illustrating a preferred embodiment of an adaptable cache of the present invention;

[00054] Fig. 6B is a block diagram illustrating components of the adaptable cache of Fig. 6A in more detail;

30 [00055] Fig. 7 is a block diagram illustrating a preferred embodiment of the present invention in which an adaptable cache is connected to a host-side bus;

[00056] Fig. 8 is a flow diagram illustrating operation of the preferred embodiment of Fig. 7;

[00057] Fig. 9 is a flow diagram illustrating proactive caching in a preferred embodiment of the present invention;

5 [00058] Fig. 10 is a flow diagram illustrating the use of data that has been proactively cached in a preferred embodiment of the present invention;

[00059] Fig. 11 is a block diagram illustrating a preferred embodiment of the present invention in which an adaptable cache is integrated with a network interface;

[00060] Figs. 12 is a flow diagram illustrating the operation of the preferred
10 embodiment of Fig. 11;

[00061] Fig. 13 is a block diagram illustrating a preferred embodiment of the present invention in which an adaptable cache is integrated with a host bus adapter;

[00062] Figs. 14 is a flow diagram illustrating the operation of the preferred embodiment of Fig. 13;

15 [00063] Fig. 15 is a block diagram illustrating a preferred embodiment of the present invention in which an adaptable cache is connected to a storage-system bus; and

[00064] Fig. 16 is a flow diagram illustrating the operation of the preferred embodiment of Fig. 15.

DETAILED DESCRIPTION OF THE DRAWINGS

20 [00065] Fig. 6A illustrates a preferred embodiment of an adaptable cache 600. As shown in Fig. 6A, adaptable cache 600 preferably comprises a data interface 610, core logic 620, and storage medium 630.

[00066] Data interface 610 preferably comprises one or more physical interface
25 connectors for connecting with one or more I/O buses, such as I/O buses 106A, B. Each connector is typically designed to connect only with a single type of I/O bus.

[00067] Core logic 620 preferably includes one or more physical components which, in the aggregate, comprise a single logical block of operation. These components preferably include:

- 30 1. bootstrapping instructions (to bring the adaptable cache from a "power off" state to an operational state);
2. an operational instruction set that defines adaptable cache functionality;
3. an interface to the storage medium; and

4. an interface to the data interface.

[00068] Storage medium 630 is adapted to store data such as digital media assets in a persistent manner. In a preferred embodiment, storage medium 630 comprises a plurality of volatile or non-volatile memory components such as DRAM or Flash memory boards that
5 utilize an electronic storage medium. Alternatively or in addition, storage medium 630 may comprise magnetic media. In a preferred embodiment, where storage medium 630 comprises volatile components, a power supply or magnetic backing may be provided to allow data to persist as if the components were non-volatile.

[00069] Electronic media does not require mechanical movement to function. The
10 absence of moving parts increases the stability of storage medium 530, decreasing the likelihood of a break down and increasing reliability. In addition, electronic media provide substantially improved data transfer capacity and improved performance since there is no mechanical seek and rotation time. As a result, electronic media are orders of magnitude faster than a conventional hard disk drive.

[00070] Aspects of core logic 620 are shown in more detail in Fig. 6B. As shown in
15 Fig. 6B, core logic 620 preferably comprises a plurality of programmable logic devices, PLDs A-C. PLD's suitable for use in the present system and method include field-programmable gate arrays ("FPGA"), complex PLDs ("CPLDs"), or other chip structures utilizing permanent, programmable, or volatile memory. Alternatively, core logic 620 may
20 comprise a plurality of CPUs, microcontrollers, or ASICs.

[00071] As further shown in Fig. 6B, PLD 622A preferably is programmed to
comprise an I/O bus controller 624 for controlling communications between adaptable cache 600 and an I/O bus to which the adaptable cache connects. Alternatively, I/O bus controller 624 may be integrated with data interface 610. PLD 622B preferably comprises bootstrap
25 instructions for taking the adaptable cache from a "power off" state to an operational state and control instructions 626 for control operation of the adaptable cache, as described in more detail below. PLD 622C is preferably programmed to provide interfaces between PLD 622A and PLD 622B, between PLD 622A and storage medium 630, and between PLD 622B and storage medium 630. Alternatively, core logic 620 may comprise one or more PLDs adapted
30 to perform all of the functions described above.

[00072] One preferred embodiment of a media server 700 incorporating an adaptable cache is shown in Fig. 7. As will be recognized, many of the components in Fig. 7

correspond to those described above in connection with Fig. 1, and corresponding components in the two figures are identified by corresponding reference numerals.

[00073] As shown in Fig. 7, adaptable cache 600 is connected to host side I/O bus 106B via a high-speed interface connection 710. In this embodiment, the physical specifications of the adaptable cache preferably include the following:

- the form factor of a PCI card;
- storage capacity in excess of 1 gigabyte (GB) using replaceable commercially off-the-shelf memory modules (such as dual inline memory modules - DIMMs) or fixed memory circuits; and
- conformity to PCI hot-swap specifications to allow the adaptable cache to be removed from service while the host system is in operation. The storage size of the adaptable cache can therefore be altered through a hot-swap without disrupting the operation of the media server.

[00074] In this embodiment, the adaptable cache and any persisted data would reside on the same I/O bus as the network interface 130, I/O bus 106B in the embodiment of Fig. 7. Data requested by the user that persists on the adaptable cache is returned via I/O bus 106B, thus reducing bus traversals.

[00075] One preferred embodiment for operation of the system shown in Fig. 7 will now be described in connection with Fig. 8. In step 801, an asset request is received at network interface 130. As noted above, this request may originate with a client, a system administrator, or another application, such as an application running on a second digital media server belonging to the same load-balancing group as server 700.

[00076] In step 802, the request is forwarded to host processor 120. Host processor 120 formulates a storage access request for the requested asset and transmits the request to adaptable cache 600 (step 803).

[00077] In step 804, adaptable cache 600 determines if the requested asset persists on its storage medium. If the asset is available, the adaptable cache returns it to host processor 120 via I/O bus 106B (step 805).

[00078] Otherwise, in step 806, the adaptable cache forwards the asset request to storage subsystem 102 for regular processing, as described below. In step 807, storage subsystem 102 returns the requested asset to host processor 120. Once it receives the asset, host processor 120 converts it to an appropriate wire format for delivery to the customer via network interface 130 (step 808).

[00079] It should be noted that although, for ease of illustration, media server 700 is described as comprising a single network interface, media server 700 may be provided with separate network interfaces for receiving client requests and transmitting digital media as taught in U.S. patent application serial No. 10/369,305, filed February 19, 2003, entitled Hybrid Streaming Platform, which is hereby incorporated by reference in its entirety for each of its teachings and embodiments. Further, it should be noted that although, for purposes of illustration, single instances of host processor 120, controller 128, host I/O bus 106B, and storage system I/O bus 106A are described, alternate embodiments of the present system and method may comprise additional numbers of these components.

[00080] In a preferred embodiment, adaptable cache 600 is adapted to proactively cache resources, and is further adapted to notify potential calling applications and other processes of assets it maintains.

[00081] Alternatively or in addition, the adaptable cache may be adapted to direct the storage system not to respond to requests for particular assets when the assets are cached in the adaptable cache. Operation of one preferred embodiment for implementing proactive caching and notification is described in connection with Fig. 9.

[00082] As shown in Fig. 9, in step 901, adaptable cache 600 monitors I/O bus 106B for asset requests. These may represent requests for content to be delivered immediately or requests for content to be delivered at a specified later time.

[00083] When a request is detected, the adaptable cache determines whether a copy of some or all of the asset is stored in storage medium 630 (step 902). In step 903, adaptable cache 600 further evaluates the request in accordance with one or more caching rules programmed into core logic 620. In a preferred embodiment, these caching rules may take account of parameters maintained by core logic 620, such as available capacity in adaptable cache 600 and the request frequency for the requested asset.

[00084] On the basis of steps 902-903, adaptable cache 600 determines whether or not some or all of the requested asset or some related asset should be proactively cached (step 904). If it is determined that some or all of an asset should be proactively cached, the system proceeds to step 905 where the adaptable cache communicates directly with the appropriate storage system or device (e.g., storage system 102) and transfers all or a portion of the asset into its storage medium 630.

[00085] In step 906, the adaptable cache notifies requesting applications and other processes that may require the requested asset of its updated content so that future requests for that asset may be directed to the adaptable cache. These applications/processes, or

associated hardware or software may preferably maintain a table that lists assets available from adaptable cache 600. Each entity receiving notification from adaptable cache 600 preferably updates its table appropriately to reflect the current content of adaptable cache 600. Processing then proceeds to step 907, described below.

5 **[00086]** If in step 904 it is determined not to cache requested content, the system proceeds directly to step 907 where parameters maintained by core logic 620 are updated. In a preferred embodiment, such parameters may, for example, include the number of times a particular asset has been requested within a specified amount of time and available capacity within the adaptable cache. Processing then returns to step 901 where adaptable cache 600
10 continues to monitor the I/O bus.

[00087] As will be recognized by those skilled in the art, passive monitoring of bus 106B by adaptable cache 600 as described above may be impractical with more modern busses which are often segmented and behave more like networks in which each device sees only traffic specifically addressed to it. Accordingly, in systems comprising such busses,
15 network interface 130 may be adapted to address each received asset request to both host processor 120 and to adaptable cache 600 so that adaptable cache 600 may monitor traffic between network interface 130 and host processor 120. References to monitoring by the adaptable cache herein should be understood to include both passive monitoring as well as monitoring using such a dual addressing scheme.

20 **[00088]** Alternatively or in addition, adaptable cache 600 may be adapted to perform interval caching wherein a sorted list of pairs of overlapping requests for the same asset is maintained that identifies pairs of requests with the shortest intervals between their start times. For these pairs, as the first request in the pair is streamed, the streamed content is also cached and then read from cache to serve the second request.

25 **[00089]** One preferred embodiment for operation of a media server 700 comprising an adaptable cache 600 adapted for proactive caching and notification will now be described in connection with Fig. 10. As shown in Fig. 10, in step 1001, a request for an asset is received via network interface 130 and forwarded to host processor 120 via I/O bus 106B. In step 1002, adaptable cache 600 monitors I/O bus 106B for such requests, caches appropriate
30 content if warranted under its caching rules, and notifies any requesting applications (including the requesting application running on host processor 120) of its updated content, as described above in connection with Fig. 9.

[00090] In step 1003, host processor 120 determines whether or not the requested asset is available from adaptable cache 600, such as by consulting a table that stores current assets

maintained by the adaptable cache. If the asset (or some portion of the asset) is available from adaptable cache 600, host processor 120 formulates a request for the asset (or portion thereof) to adaptable cache 600 (step 1004). In step 1005, adaptable cache 600 returns the requested asset to host processor 120.

5 **[00091]** Otherwise, if the asset is not available from adaptable cache 600, host processor 120 formulates a request for the asset to storage system 102 (step 1006). The requested asset is read in blocks from a storage device 104 of storage system 102 and transmitted to host processor 120, as shown by the iteration of steps 1007-1010. More particularly, for each block, the storage device finds the block on the hard drive (step 1007),
10 reads the block (step 1008), transmits the block (step 1009), and determines whether or not the asset comprises additional blocks (step 1010).

[00092] Another preferred embodiment for implementing the present system and method is shown in connection with Fig. 11. As in Fig. 10, adaptable cache 600 in Fig. 11 also resides as a device connected to the host side I/O bus 106B. In this embodiment,
15 however, adaptable cache 600 is preferably integrated with network interface 130. The adaptable cache 600 preferably interconnects with host side I/O bus 106B via interface connection 146. Preferred physical specifications for the adaptable cache in this preferred embodiment comprise:

- the form factor of a network interface card (e.g., a peripheral component
20 interconnect or PCI card) which may be plugged into an available expansion slot on the host system (e.g., a PCI slot);
- storage capacity in excess of 1 gigabyte (GB) using replaceable commercially off-the-shelf memory modules (such as dual inline memory modules - DIMMs) or fixed memory circuits; and
- 25 • conformity to PCI hot-swap specifications to allow the adaptable cache to be removed from service while the host system is in operation. As noted above, the storage size of the adaptable cache can therefore be altered through a hot-swap without disrupting the operation of the media server.

[00093] In this preferred embodiment, adaptable cache 600 is programmed to respond
30 directly to asset requests when the requested asset is available in its storage medium 630. In this way, asset requests may be serviced and delivered from the network interface card, eliminating bus traversals when assets requested by the user reside in the adaptable cache.

[00094] Operation of the system shown in Fig. 11 will now be described in connection with Fig. 12. In step 1201, an asset request is received at network interface 130. In step 1202, adaptable cache 600 determines if the requested asset is available on the adaptable cache.

5 **[00095]** If the asset is available on the adaptable cache, the request is preferably serviced and delivered to the user from the same card, eliminating bus traversals on buses 106 (step 1203). More specifically, the adaptable cache retrieves the resource from storage medium 630, converts it to an appropriate wire format and delivers it to the requesting client.

[00096] Otherwise, in step 1204, if the requested resource is not available from the adaptable cache, the request is forwarded to host processor 120 for processing. In step 1205, host processor 120 formulates a request for the asset to storage system 102. In step 1206, the asset is returned to host processor 120, as described above in connection with Fig. 10. In step 1207, host processor 120 converts the asset to an appropriate wire format and delivers it to the client via network interface 130.

15 **[00097]** It should be recognized that the proactive caching and notification described above may also be implemented in this embodiment. Thus, adaptable cache 600 may be adapted to monitor received requests, proactively cache some or all of an asset in accordance with caching rules, and notify one or more applications or processes of content that it is currently storing. Further, the adaptable cache may be adapted to direct the storage system not to respond to requests for particular assets when the assets are cached in the adaptable cache.

[00098] Another preferred embodiment for implementing the present system and method is shown in Fig. 13. In the embodiment of Fig. 13, adaptable cache 600 is integrated with controller 128, bridging I/O buses 106A, B. In this embodiment, adaptable cache 600 preferably plugs into an expansion slot on the host system and provides multiple standard high-speed interfaces, such as bridging Fibre Channel and PCI I/O interfaces. In this embodiment, preferred physical specifications of the adaptable cache include:

- the form factor of a peripheral component interconnect (PCI) card;
- storage capacity in excess of 1 gigabyte (GB) using replaceable commercially off-the-shelf memory modules (such as dual inline memory modules - DIMMs) or fixed memory circuits; and
- conformity to PCI hot-swap specifications to allow the adaptable cache to be removed from service while the host system is in operation. As noted above, the

storage size of the adaptable cache can therefore be altered through a hot-swap without disrupting the operation of the media server.

[00099] Operation of the system shown in Fig. 13 will now be described in connection with Fig. 14. In step 1401, a user request is received at network interface 130. In step 1402, the request is forwarded to host processor 120 via I/O bus 106B. In step 1403, host processor 120 sends a request for the asset to storage system 102 via I/O bus 106B.

[000100] In step 1404, adaptable cache 600 (integrated with controller 128 in this embodiment) monitors asset requests that traverse I/O buses 106A, B and determines if the requested asset is available on the adaptable cache. In step 1405, if the asset is available on the adaptable cache, it is returned to host processor 120.

[000101] Otherwise, if the requested resource is unavailable from the adaptable cache, the request is forwarded to storage system I/O bus 106A for delivery to the appropriate storage device 104 where the resource persists (step 1406). In step 1407, the storage device returns the resource to the requesting application, as described in more detail above. In step 1408, host processor 120 receives the requested resource, as described in more detail above.

[000102] It should be recognized that the proactive caching and notification described above may also be implemented in this embodiment. Thus, adaptable cache 600 may be adapted to monitor received requests, proactively cache some or all of an asset in accordance with caching rules, and notify one or more applications or processes of content that it is currently storing. Further, the adaptable cache may be adapted to direct the storage system not to respond to requests for particular assets when the assets are cached in the adaptable cache.

[000103] Yet another preferred embodiment for implementing the present system and method is shown in Fig. 15. In Fig. 15, an adaptable cache 600 resides on storage system 102. Adaptable cache 600 preferably interconnects with storage system I/O bus 106A via a high-speed interface connection 1500. This high-speed interface connection preferably allows adaptable cache 600 to replace or supplement existing hard drive storage devices on storage system 102 (including RAID arrays or JBODs) as long as the system has a compatible receptacle and I/O interface. In this embodiment, preferred physical specifications of the adaptable cache comprise:

- the form factor of a 3.5" hard disk drive with a 1" nominal height;
- dual fibre channel interface utilizing a standard SCA 40-pin connector and operating at transfer rates of either 1 or 2 gigabits per second (Gbps);

- storage capacity in excess of 1 gigabyte (GB) using replaceable commercially off-the-shelf memory modules (such as dual inline memory modules - DIMMs) or fixed memory circuits - this facilitates a lower cost while simultaneously providing the benefit of readily available and quality controlled components; and

- hot-swap capability (the ability to swap or remove the adaptable cache from service while the system is in operation). As noted above, the storage size of the adaptable cache can therefore be altered through a hot-swap without disrupting the operation of the media server.

[000104] Operation of the preferred embodiment shown in Fig. 15 will now be described in connection with Fig. 16. In step 1601, a user request is received at network interface 130. In step 1602, the request is forwarded to host processor 120 via I/O bus 106B. In step 1603, host processor 120 sends a request for the asset to storage system 102 via I/O bus 106B.

[000105] In step 1604, adaptable cache 600 monitors asset requests that traverse I/O bus 106A and determines if the requested asset is available on the adaptable cache. As noted above, those skilled in the art will recognize that passive monitoring of bus 106B by adaptable cache 600 may be impractical with more modern busses which are often segmented and behave more like networks in which each device sees only traffic specifically addressed to it. Accordingly, as noted above, in systems comprising such busses, host processor 120 may be adapted to address each received asset request to both storage device 104 and to adaptable cache 600 so that adaptable cache 600 may monitor traffic between host processor 120 and storage device 104.

[000106] In step 1605, if the asset is available on the adaptable cache, it is returned to host processor 120. In this case, the adaptable cache or other suitable component in storage system 102 may also preferably be adapted to preclude other storage devices 104 from responding to the request from host processor 120 since such storage device will be unable to retrieve and forward the asset to host processor 120 as efficiently as adaptable cache 600 (step 1606).

[000107] Otherwise, if the requested resource is unavailable from the adaptable cache, the request is delivered to the appropriate storage device 104 where the resource persists (step 1607). In step 1608, the storage device returns the resource to the requesting application, as described in more detail above.

[000108] It should be recognized that the proactive caching and notification described above may also be implemented in this embodiment. Thus, adaptable cache 600 may be adapted to monitor received requests transmitted via I/O bus 106A, proactively cache some or all of an asset in accordance with caching rules, and notify one or more applications or processes of content that it is currently storing. Alternatively, these caching and monitoring components may be divided. More specifically, a separate monitoring component may be provided on I/O bus 106A to monitor requests as they are received by network interface 130. When appropriate, the monitoring component may instruct adaptable cache 600 (residing, for example, on I/O bus 106A) to retrieve and store some or all of an asset.

[000109] It should also be noted that although, in the preferred embodiments described above, system components are linked via PCI buses such as bus 106A, B, these components may alternatively be linked via other bus types or data exchanges such as switched fabric and associated daughtercards.

[000110] While the invention has been described in connection with specific

embodiments, it is evident that numerous alternatives, modifications, and variations will be apparent to those skilled in the art in light of the foregoing description.